VERITAS® Cluster Server

Quick-Start Guide

An Example with NFS

Version 1.1

Solaris

August 1999

P/N 100-001205

TRADEMARKS

Table of Contents

Preface

Welcome to VERITAS® Cluster Server™ (VCS), the latest high-availability solution for Solaris clusters. VCS enables you to monitor systems and application services, and to restart services on a different system when hardware or software fails. Its enhanced feature set offers:

- *Increased scalability.*

  VCS provides a framework that supports parallel applications as they distribute load on multiple systems. It monitors the application services, and can be configured to fail over faulted instances to standby systems.

- *Enhanced shared storage.*

  VCS supports shared storage on clusters of up to 32 systems connected by any combination of systems, disks, and other devices. VCS is designed to grow with your organization, enabling you to migrate from shared SCSI clusters to fibre channel and other SAN technologies.

- *Simplified administration.*

  VCS's replication technology provides a single-system view of administration. Systems share identical copies of configuration and resource states, enabling you to monitor and control applications and reconfigure the cluster from any system. Using the Java-based graphical user interface, you can administer multiple clusters from any Solaris or Windows NT workstation. VCS even supports online addition of new systems: you can add systems without taking down applications or the VCS services that monitor them.

How This Guide Is Organized

Chapter 1, “Installing VCS,” instructs you on how to install and configure your hardware on two systems, and how to install the VCS software.

Chapter 2, “Configuring VCS,” introduces the configuration wizard and explains what is required to run it.

Chapter 3, “Basic Operations,” describes how to perform primary tasks.

Chapter 4, “Troubleshooting VCS,” provides a list of potential errors and ways to correct them.

Technical Support

For assistance with any of the VERITAS products, contact our Customer Support hotline:

U.S. and Canadian Customers: 1-800-342-0652

International Customers: +1 (650) 335-8555

Fax: (650) 335-8428

Email: support@veritas.com

Conventions

Typeface	Usage
courier	computer output, files, attribute names, device names, and directories
courier (bold)	user input and commands, keywords in grammar syntax
italic	new terms, titles, emphasis, variables replaced with a name or value
italic (bold)	variables within a command
Symbol	**Usage**
%	C shell prompt
$	Bourne/Korn shell prompt
#	Superuser prompt (for all shells)
\	Command-line continuation if last character in line. Not to be confused with an escape character.

Installing VCS

1

Introduction

This chapter describes how to install VCS on a two-system cluster. Specifically, you will be instructed on how to employ VCS in a *symmetric* failover configuration to provide highly available NFS service. In a symmetric configuration, if one system fails, its partner assumes control of the services while continuing to maintain its own. The examples in this guide support symmetric configurations using disk partitions and the VERITAS® Volume Manager™.

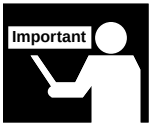

This exercise is designed to familiarize new VCS users with the product's basic installation and configuration procedures, and with its primary operations. To configure multiple systems, or to explore other configuration options, refer to the *VERITAS Cluster Server User's Guide.*

Requirements

Item	Description
VCS systems	Two SPARC systems running Solaris 2.5.1 or later.
CD-ROM drive	One CD-ROM drive on each system, or a drive accessible to both.
Disk partitions	This configuration requires shared disk storage for file systems being exported for NFS service.
Ethernet controllers	This configuration requires two Ethernet interfaces per system, in addition to the built-in public Ethernet interface.
SCSI adapters	VCS requires one built-in SCSI adapter per system for the operating system disks, and one or more additional SCSI adapters per system for the shared data disks.
Disk space	Each VCS system must have at least 35 megabytes of free space in the `/opt` file system.
RAM	Each VCS system requires at least 128 megabytes.
IP addresses	Each group of services exported from the cluster must have a unique IP address. **Note:** A typical dual-system failover configuration (symmetric) requires four IP addresses: two for the systems, and two for the highly available services that move between systems.

Overview

Installing VCS involves various tasks, and each task must be performed in the order presented below. Instructions begin on page 4.

- ✓ Setting up your hardware.
- ✓ Identifying the block devices for NFS services.
- ✓ Checking the major and minor numbers.
- ✓ Preparing the NFS services.
- ✓ Installing the VCS software.
- ✓ Establishing the cluster.

Setting Up the Hardware

VCS requires specific connections from each system to the shared disks, and from each system to the public and private networks. Review the figure below.

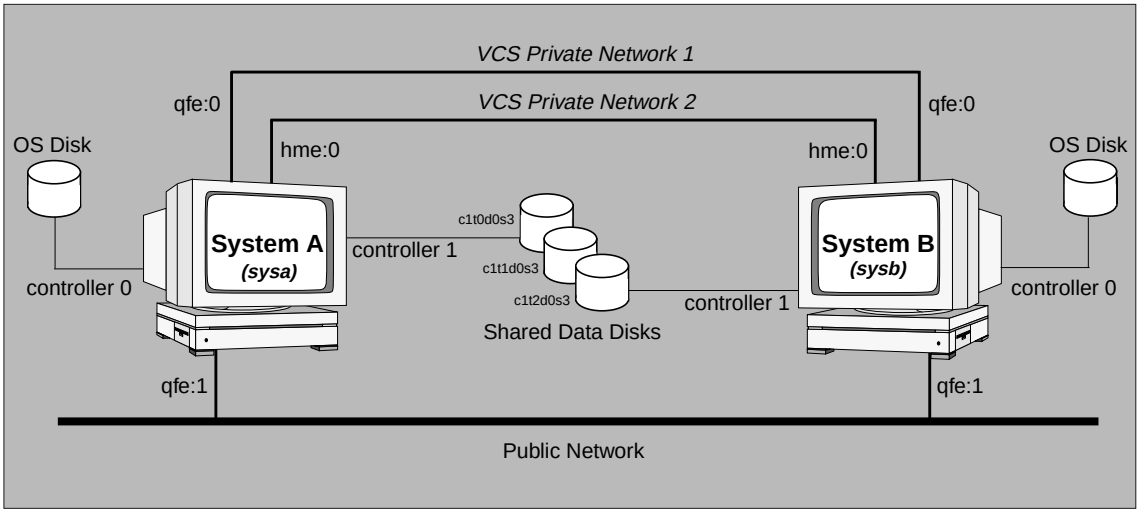

Example of a VCS Two-System Configuration

To set up the hardware:

1. Install the required Ethernet and SCSI controllers.
2. Connect the VCS private Ethernet controllers on each system. This can be done with cross-over Ethernet cables, or with hubs that are separate from the hubs used for the public network. If you're using hubs, use one for each VCS private network.
3. Cable the shared disks to System A and terminate the SCSI bus.

4. From the EEPROM prompt (ok) on System A, probe your controllers and devices to confirm you have installed them correctly. *Devices on the shared SCSI bus must not conflict with System A's SCSI ID 7 or System B's SCSI ID 5.* Type the following commands:

   ```
   ok show-devs
   ok probe-scsi-all
   ```

5. Now cable the shared disks between the two systems, as illustrated below:

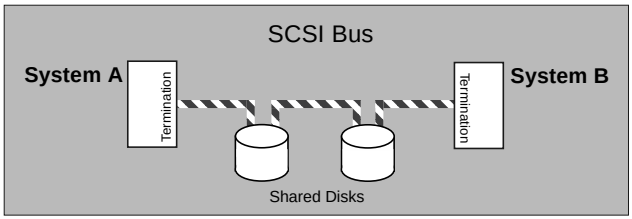

6. Modify the SCSI ID on System B. This ensures that each system has a different ID number. If you have more than one SCSI adapter on the system, this number cannot be used on any of the buses attached to System B.

   From the EEPROM prompt on System B, type:

   ```
   ok setenv scsi-initiator-id 5
   ```

7. Boot each system:

   ```
   ok boot -r
   ```

Configuring the Maintenance IP Address

The term “maintenance IP address” refers to the IP address assigned to each system during Solaris installation. It is a public network connection initialized by Solaris at system boot and is not under control of VCS.

The recommended hardware configuration for VCS may require that you move the maintenance IP address from one interface to another. For example, if your network interfaces include an interface built into the motherboard of your system (`le0`), and four interfaces on an expansion card (`qfe0`, `qfe1`, `qfe2`, and `qfe3`), VCS requires that you move the maintenance IP address from `le0` to a `qfe` interface. The purpose of moving the maintenance IP (and the public network connection) is to create a private network consisting of two or more independent hardware components. In this example, you would configure the maintenance IP on `qfe1`, and the VCS private network on `qfe0` and `le0`, thereby protecting the cluster against network partitions caused by the failure of a single component, such as the qfe card.

The following figure demonstrates the maintenance IP configuration for sysa and sysb:

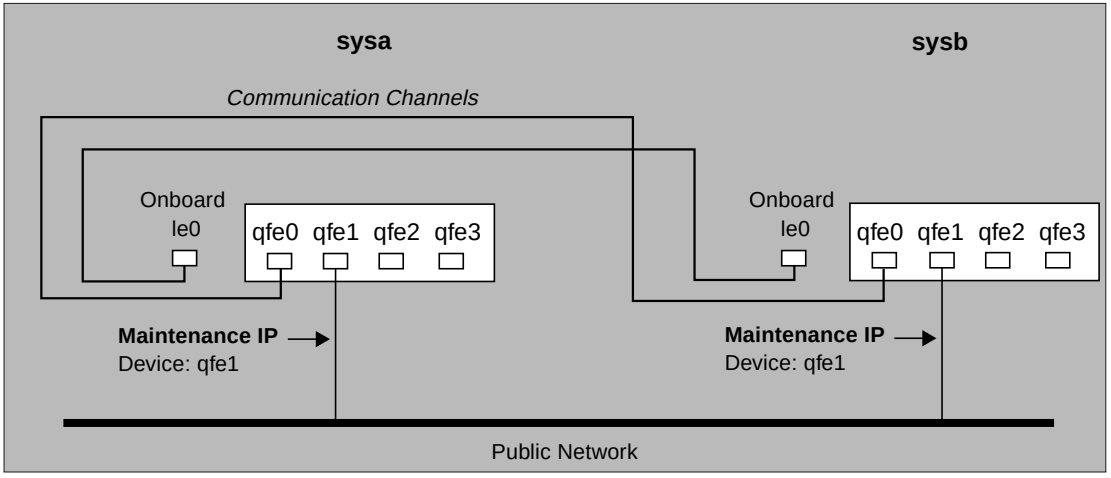

To move the maintenance IP, perform the following steps on each system. Note that the following procedure is an example of moving the maintenance IP address on a system from `le0` to `qfe1`. Your network interface names may differ.

1. Confirm that the public network is connected by pinging a known machine that is external to the cluster.
2. Confirm the configuration of the maintenance IP on `le0`:

   ```
   # cat /etc/hostname.le0
   ```

   Output displays the name of the system (in this example, sysa).
3. Rename this file to reflect the new network interface:

   ```
   # mv /etc/hostname.le0 /etc/hostname.qfe1
   ```
4. Unplug the public service network cable from the built-in network interface and plug it into the Quad-Fast-Ethernet card in plug number 1 (second from left).
5. Reboot the system.
6. Confirm that the public network is still connected by pinging a known machine that is external to the cluster.

Note: The device used for the maintenance IP device should not be used as a network communication channel. However, it can be shared with service groups to provide network and IP resources.

Network Partitions and the Sun Boot Monitor

Sun SPARC systems provide a console-abort sequence that enables you to halt and continue the processor: `L1-A` or `STOP-A` on the keyboard, `BREAK` on the serial console input device. Each command is then followed by a response of `go` at the `ok` prompt. Continuing operations after the processor has stopped may corrupt data and is therefore unsupported by VCS. Specifically, when a system is halted with the abort sequence it stops producing heartbeats. The other systems in the cluster then consider the system failed and take over its services. If the system is later enabled with `go`, it continues writing to shared storage as before, even though its applications have been restarted on other systems.

Solaris 2.6

In Solaris 2.6, Sun introduced support for disabling the console-abort sequence. We recommend disabling the sequence on systems running Solaris 2.6 or greater. To disable the console-abort sequence:

1. Add the following line in the file `/etc/default/kbd` (create the file if it does not exist):

   ```
   KEYBOARD_ABORT=disable
   ```

2. Reboot.

3. If necessary, see the manual page `kbd(1)` for details.

Solaris 2.5.1

If you do not want to disable the abort sequence, do not type `go` at the `EEPROM ok` prompt. If a system has been stopped with the abort sequence in a VCS cluster, type `boot` at the `ok` prompt.

Identifying Block Devices for NFS Service

Your configuration includes disks on the shared bus that support NFS service. For symmetric failover configurations these disks are organized into two groups: one group per system, upon which the NFS file systems are deported during normal operation.

Each group can support multiple file systems, and each file system can be configured on a disk partition or on a VERITAS Volume Manager volume. An example disk partition name is `/dev/dsk/c1t1d0s3`. An example volume name is `/dev/vx/dsk/shareydg/vol3`. Each name represents the block device on which the file system will be mounted.

For Volume Manager configurations, install the Volume Manager as instructed in the *VERITAS Volume Manager Installation Guide*. Disks on the shared bus must be configured into disk groups other than `rootdg`. Create NFS service disk groups on one system only. The disk groups will be exported and imported onto the other system by VCS.

Checking Major and Minor Numbers

Block devices providing NFS service must have the same major and minor numbers on each system. Major and minor numbers are used by Solaris to identify the logical partition or disk "slice." NFS also uses them to identify the exported file system. These numbers must be checked to ensure that the NFS identity for the file system is the same when exported from each system.

To check major and minor numbers:

1. Use the following command on both systems to display the major and minor numbers for a block device. For Volume Manager volumes, you must first import the associated shared disk group on each system, one system at a time.

   # **ls -lL** *block_device*

   The variable *block_device* refers to a partition on which a file system is mounted for export via NFS. This procedure must be followed for each NFS file system.

   For example, type:

   # **ls -lL /dev/dsk/c1t1d0s3**

   Output on System A resembles:

   ```
   crw-r-----  1 root  sys  32,134 Dec 3 11:50 /dev/dsk/c1t1d0s3
   ```

   Output on System B resembles:

   ```
   crw-r-----  1 root  sys  32,134 Dec 3 11:55 /dev/dsk/c1t1d0s3
   ```

   Notice that the major and minor numbers shown above match; both major numbers are 32, and both minor numbers are 134.

2. If either the major or minor numbers do not match, see "Reconciling Major / Minor Numbers" on page 39.

3. Repeat steps 1–2 for each block device used for NFS.

Preparing NFS Services

Before configuring the VCS software, you must initialize the disks, file systems, and mount points in preparation for your NFS services.

For each of the file systems that will support NFS service, complete the steps below:

1. Create the file system using `mkfs(1M)` on one system only.
2. Create the mount points on both systems using `mkdir(1M)` on both systems.

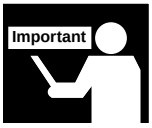

Do not add exported file systems to `/etc/vfstab` or `/etc/dfs/dfstab`. VCS will mount and export these file systems automatically.

Installing the VCS Software

To mount the VERITAS CD-ROM on both systems:

1. Insert the CD into a drive connected to each system.
   - If you are running Solaris volume-management software, the software automatically mounts the CD as `/cdrom/cdrom0`.
   - If you are not running Solaris volume-management software, you must mount the CD manually. For example:

     ```
     # mount -F hsfs -o ro /dev/dsk/c0t6d0s2 \
           /cdrom/cdrom0
     ```

     Note that `/dev/dsk/c0t6d0s2` is the default name for the CD drive.

2. Type the following command on both systems to install the VCS software:

   ```
   # pkgadd -d /cdrom/cdrom0
   ```

   You will receive a message listing the available packages.

3. When prompted, select `all`.
4. As the packages are being added, answer `Yes` when prompted.
5. After all the packages are added, type `q`. The following message is displayed:

   ```
   *** IMPORTANT NOTICE ***
          This machine must now be rebooted in order to
          ensure sane operation. Execute
                shutdown -y -i6 -g0
          and wait for the "Console Login:" prompt.
   ```

CAUTION! Do not reboot at this time. Before rebooting, you must establish the cluster. Proceed to the next section, “Establishing the Cluster” on page 13.”

Establishing the Cluster

VCS provides a wizard named "Quick-Start" to help you establish the cluster. Quick-Start is run on a single VCS system, the *local* system, and configures the local system and its partner, the *remote* system. Instructions on how to use Quick-Start begin on page 16.

Note: This wizard will assist you in establishing a two-system cluster connected by dual Ethernet heartbeat links (communication channels). Quick-Start applies specifically to the sample configuration presented in this guide and is not applicable to other configurations.

To establish your cluster:

1. Review the definitions below.
2. Complete the worksheet on page 15.
3. Run the VCS Quick-Start wizard.
4. Verify LLT, GAB, and the cluster.

Remote System Name

An IP name or number of the remote VCS system; for example, sysb. To use the Quick-Start wizard, the local VCS system must have `rsh(1M)` access as `root` on the remote system.

Cluster Name

The name assigned to your cluster. The default name is vcs.

System List

A list of all systems in the cluster; for example, sysa and sysb. These names must correspond to the values returned by the `uname(2)` system call. These are also the names used by the VCS graphical user interface. Note that the VCS Quick-Start wizard provides defaults based on `uname(1)`.

Communication Channels

VCS requires two or more communication channels between systems in the cluster. In the sample configuration presented in this guide, these channels consist of two private networks, which are separate from the public network used by clients for NFS service. (Review the figure on page 4.) VCS uses its own protocol on the private networks. Do not configure IP on these networks unless you are using the IP for a purpose other than VCS. Specifically, do not create an `/etc/hostname.`*device* file for these interfaces.

For example, channel information for a two-system cluster would resemble:

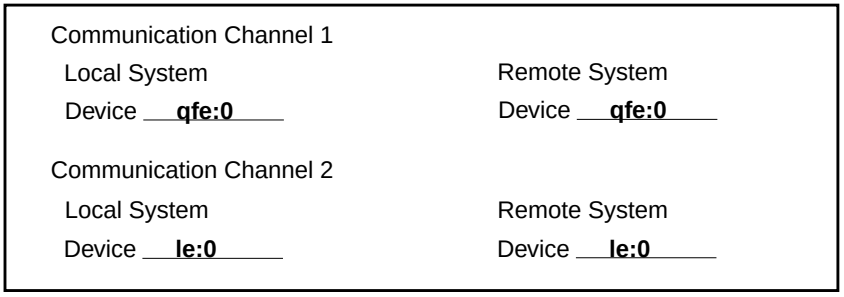

The Quick-Start Worksheet

Complete the following worksheet to ensure that you provide the correct information to the Quick-Start wizard.

Remote System Name _____

Cluster Name _____

System List

Local System _____

Remote System _____

Communication Channel 1

Local System	Remote System
Device _____	Device _____

Communication Channel 2

Local System	Remote System
Device _____	Device _____

Running the Quick-Start Wizard

The Quick-Start wizard is run from a single VCS system. To establish your cluster, the wizard will prompt you for the cluster name, system list, and communication channels, then create the following configuration files on the local and remote systems:

- `/etc/llthosts`
- `/etc/lltab`
- `/etc/gabtab`
- `/etc/VRTSvcs/conf/config/main.cf`
- `/etc/VRTSvcs/conf/config/sysname`

Note that the system running the wizard is referred to as the local system, and the second system in the cluster is referred to as the remote system.

Before running the Quick Start, make sure you've met the following requirements:

✓ You must be `root`.

✓ The local system has `root rsh(1M)` access to the remote system.

✓ You have set up your hardware as instructed on page 4.

✓ You have installed the VCS software as instructed on page 12.

Note: If you make a mistake while entering your input, type `Control-B` to return to the previous screen. To exit the wizard, type `Control-C`. (Note that if you exit the wizard before establishing your cluster, you must start again from the beginning.)

To run the VCS Quick-Start wizard:

1. Type the following command:

   ```
   # $VCS_HOME/wizards/config/quick_start
   ```

   If it returns an error message about screen size, adjust your screen size using the `stty` variable. Set the minimum screen rows to 24 and the minimum screen cols (columns) to 80.

2. Refer to the information you entered in the configuration worksheet (page 15) to answer subsequent prompts.

3. You will receive a confirmation message after the wizard successfully runs the quick start:

   ```
   Congratulations! You have completed the initial VERITAS
   Cluster Server installation and configuration.
   ```

Verifying LLT, GAB, and the Cluster

After running the VCS Quick-Start wizard, you must verify that LLT, GAB, and the cluster are up and running.

To verify that LLT and GAB are up and running:

Type the following command as `root` on each system:

```
# /sbin/gabconfig -a
```

If LLT and GAB are up and running, GAB port membership information is returned:

```
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

`Port a` indicates that LLT and GAB are communicating; `gen a36e0003` is a random generation number; `membership 01` indicates that systems `0` and `1` are connected.

`Port h` indicates that VCS is started; `gen fd570002` is a random generation number; `membership 01` indicates that systems `0` and `1` are both running VCS.

If LLT and GAB are not up and running, no GAB port membership information is returned:

```
GAB Port Memberships
===================================
```

For more information on LLT and GAB, refer to the *VERITAS Cluster Server Installation Guide* and the *VERITAS Cluster Server User's Guide.*

To verify that the cluster server is up and running

Type the following command as `root` on each system:

```
# /opt/VRTSvcs/bin/hasys -display
```

If the cluster is up and running, the attributes of the active cluster are returned:

```
# $VCS_HOME/bin/hasys -display

#System  Attribute          Value
thor8    AgentsStopped      0
thor8    ConfigBlockCount   0
thor8    ConfigCheckSum     0
.
.
.
.
```

Note that additional information may be displayed.

If the cluster is not up and running, the following message is returned:

```
Unsuccessful open of service
```

For more information on the `hasys -display` command, refer to the *VERITAS Cluster Server User's Guide.*

Completing the Cluster Configuration

Proceed to Chapter 2 to complete the cluster configuration.

Configuring VCS

Overview

VCS also provides a wizard named “Quick-NFS” to help you build a sample configuration after you’ve established the cluster.

Before running Quick-NFS, complete the configuration worksheet on page 24. Keep this information available to ensure that you answer each prompt correctly.

This configuration procedure consists of:

✓ Defining VCS service groups and resources.

✓ Designating file systems for each service group.

Defining Service Groups

You will be prompted for information regarding *service groups*. In VCS, each service is specified as a group. Each service group consists of resources that must be started (brought online) as part of making the group available. A single resource can belong to only one service group.

There are many types of resources, including:

- Hardware resources, such as disks and network cards.
- Operating system resources, such as disk groups, IP addresses, file systems, and volumes.
- Programs and databases that export NFS file systems.

You will configure two service groups: “groupx” and “groupy.” The local system (the system running the Quick-NFS wizard) is the primary system for groupx, and the remote system is the primary for groupy. The primary system is where the application automatically starts when the system is booted.

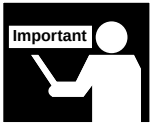

As mentioned previously, this guide instructs you on how to configure VCS for two-system symmetric failover of an NFS service. This type or “class” of configuration requires two service groups. Each service group can support multiple file systems (examples in this guide illustrate two per group), and each file system can be configured on disks or on volumes using the VERITAS Volume Manager.

Service Group Name

A VCS service group is an abstract container of related resources that can be administered as a unit. For example, the service group name is used to bring the group online, take the group offline, or to migrate the service from one system in the cluster to another.

Public Network Device Name

The network device on which the application services run.

Service Group IP Address

The IP address of the hostname described above. Clients use this address to attach to the system running the service group.

Do not create an `/etc/hostname.`*device* for this address. This IP address is configured by VCS when the service is brought online, not during the system's boot sequence.

For example, information regarding service groupx would resemble:

Service Group

Public Network Device Name **le0**

Service Group Name **groupx**

Service Group IP Address **192.200.40.101**

Designating File Systems

Next, you will be prompted for information regarding the *file systems* for each service group. This information is specified for each file system exported via NFS.

Note: If you are using volumes, you must specify a disk group. This item appears in the worksheet under “Service Group X” and “Service Group Y.”

Review the following definitions and sample worksheet entries.

Using VERITAS Volume Manager Volume?

Answer no to configure this file system on a disk partition. Answer yes to configure it on a VERITAS Volume Manager volume.

Disk Partition/Volume Name

Depending on what you answered above, you will be prompted for one of the following responses:

- The disk partition (slice) on which to mount the file system.
- The name of the VERITAS volume on which to mount the file system.

Mount Point

The directory on which the file system is mounted.

Mount Options

Options given to the mount command, such as “ro” for read-only or “rw” for read / write.

File System Type

The type of file system, such as ufs or vxfs.

Disk Group

The VERITAS Volume Manager disk group associated with this service group. All volumes associated with this service group are configured in the disk group. (When using the Quick-NFS wizard, you will be prompted for this information only if you have configured the Volume Manager.)

For example, information regarding a file system configured on a disk partition would resemble:

If configured on a volume:

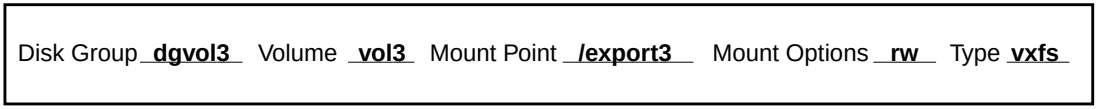

The Quick-NFS Worksheet

Complete the following worksheet before running the Quick-NFS wizard. If necessary, refer to the previous sample entries for assistance.

Service Groupx

Public Network Device Name _____

Service Group Name _____

Service Group IP Address _____

Disk Group *(optional)* _____

Disk Partition/Volume	Mount Point	Mount Options	Type

Service Groupy

Public Network Device Name _____

Service Group Name _____

Service Group IP Address _____

Disk Group *(optional)* _____

Disk Partition/Volume	Mount Point	Mount Options	Type

Running the Quick-NFS Wizard

The Quick-NFS wizard first verifies the cluster, then creates a configuration file and propagates the file to a remote system. All configuration directories reside under `/etc/VRTSvcs/conf`.

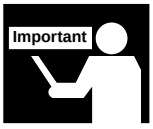

If you want to verify the cluster yourself or view the syntax of the configuration, look at the file created by the VCS Quick-Start wizard: `/etc/VRTSvcs/conf/config/main.cf`. The Quick-NFS wizard appends information to this file. We recommend that you make a backup copy of this file before you edit it or run the Quick-NFS wizard.

Note: If you make a mistake while entering your input, type `Control-B` to return to the previous screen. To exit the wizard, type `Control-C`. (Note that if you exit the wizard before building your configuration, you must start again from the beginning.)

Type the following command to run the Quick-NFS wizard:

```
# $VCS_HOME/wizards/services/quick_nfs
```

A text run of this wizard begins on the next page. Bold text indicates the examples used in the sample configuration.

```
Quick-NFS Install Wizard                                SCREEN # 1

WELCOME TO THE QUICK-NFS WIZARD!

This wizard will help you set up an NFS service on the two-system VCS
cluster that you have already configured by using the Quick-Start
wizard. You will be asked a series of questions that will be used in
configuring the NFS service.

^F=Forward screen ^C=Exit
```

```
Quick-NFS Install Wizard                                SCREEN # 2

SERVICE GROUP INFORMATION FOR GROUPX

The primary system for this service group is sysa. Select a network
device from the list shown below. Your selection will be the default
network device. To enter any other option, type in the new name.

                    [j=down/k=up]
                                  le0
                                  qfe0
                                  qfe1
                                  qfe2
                                  [MORE..]

Public Network Device Name:       le0
Service Group Name:               groupx
Service Group IP Address:         192.200.40.101

^B=Back screen ^C=Exit
```

```
Quick-NFS Install Wizard                                SCREEN # 3

You can configure a service group with disks or volumes. If you have
installed the VERITAS Volume Manager and want to configure volumes in
your service group, answer yes to the following question.

Do you want to use the VERITAS Volume Manager?:    n

Quick-NFS Install Wizard                                SCREEN # 4

CONFIGURE FILE SYSTEM FOR SERVICE GROUPX

Select a disk partition from the list shown below. You selection will
be the default disk partition. To enter any other option, type in the
new name.

                    [j=down/k=up]
                                  /dev/rdsk/c0t0d0s0
                                  /dev/rdsk/c0t0d0s1
                                  /dev/rdsk/c0t0d0s2
                                  /dev/rdsk/c1t0d0s0
                                  [MORE..]

Disk Partition:                   /dev/rdsk/c1t0d0s0
Mount Point:                      /export1
Mount Options:                    rw
File System Type:                 ufs

Additional File Systems for this group?:    n

^B=Back screen ^C=Exit
```

```
Quick-NFS Install Wizard                                SCREEN # 5

SERVICE GROUP INFORMATION FOR GROUPY

The primary system for this service group is sysb. Select a network
device from the list shown below. Your selection will be the default
network device. To enter any other option, type in the new name.

                    [j=down/k=up]
                                  le0
                                  qfe0
                                  qfe1
                                  qfe2
                                  [MORE..]

Public Network Device Name:       le0
Service Group Name:               groupy
Service Group IP Address:         192.200.40.102

^B=Back screen ^C=Exit

Quick-NFS Install Wizard                                SCREEN # 6

You can configure a service group with disks or volumes. If you have
installed the VERITAS Volume Manager and want to configure volumes in
your service group, answer yes to the following question.

Do you want to use the VERITAS Volume Manager?:    y
```

```
Quick-NFS Install Wizard                                SCREEN # 7

CONFIGURE FILE SYSTEM WITH VOLUMES FOR SERVICE GROUPY

Select a volume from the list shown below. To enter any other option,
type in the new name.

                    [j=down/k=up]
                                  rootdisk3vol
                                  rootdisk4vol
                                  rootvol
                                  swapvol

Volume Name:                      vol3
Mount Point:                      /export3
Mount Options:                    ro
File System Type:                 vxfs
Disk Group                        dgvol3

Additional File Systems for this group?   n

Note: The same disk group will be used if you configure additional file
systems.

^B=Back screen ^C=Exit
```

```
Quick-NFS Install Wizard                                SCREEN # 8

EXIT THE QUICK-NFS WIZARD

Congratulations! You have successfully created a valid NFS
configuration file.

The configuration file has been copied to
/etc/VRTSvcs/conf/config/main.cf.

Important: To run VCS you must reboot both systems.

^B=Back screen ^C=Exit
```

Basic Operations

This chapter explains basic VCS operations. For a complete list of VCS commands, or for more information regarding how VCS operates, refer to the *VERITAS Cluster Server User's Guide*.

In this chapter, you will learn how to perform these basic tasks:

- start VCS (page 32)
- verify the state of groups and resources
- start, stop, and migrate a group
- test failover
- use the graphical user interface (GUI) to monitor the cluster

Starting VCS

The VCS Quick-Start wizard started VCS for you; however, sometimes you may have to start it manually.

To begin, place the VCS command directory in your path. For example:

```
# export PATH=$PATH:/opt/VRTSvcs/bin
```

To start `groupx` and `groupy` on Systems A and B, type the following command on both systems:

```
# hastart
```

Note: As in the previous examples, this configuration consists of System A and System B, and groupx and groupy. The configuration file is stored in the directory `/etc/VRTSvcs/conf/config`. Note that after VCS is started, commands to VCS may be entered on either system.

Verifying the State of Systems and Groups

To verify that both systems are operating and both groups are online, type the following command on either system:

```
# hastatus -summary
```

Output from the command resembles the following example. (Bold text indicates that systems and groups are online.)

```
--SYSTEM STATE
--System              State              Frozen
A sysa                RUNNING            0
A sysb                RUNNING            0


--GROUP STATE
--Group      System   Probed   AutoDisabled   State
B groupx     sysa     Y        N              ONLINE
B groupy     sysb     Y        N              ONLINE
```

If the groups are not online when you enter the command, they may be in the process of coming online. Type the following command on either system to provide continual updates as states change in the cluster:

```
# hastatus
```

Note: If the groups fail to go online within two minutes, see Appendix A in the *VERITAS Cluster Server User's Guide.*

Verifying the State of Resources

To verify the state of the resources, type the following command on either system:

```
# hares -display nfs_export1
```

Output from this command resembles the example below. (Bold text indicates that the Share resource `nfs_export1` is online on `sysa`.)

```
#Resource    Attribute          System   Value
nfs_export1  ConfidenceLevel    sysa     0
nfs_export1  ConfidenceLevel    sysb     0
nfs_export1  Probed             sysa     1
nfs_export1  Probed             sysb     1
nfs_export1  State              sysa     ONLINE
nfs_export1  State              sysb     OFFLINE
nfs_export1  ArgListValues      sysa     /export1/dev/...
nfs_export1  ArgListValues      sysb     /export1 /dev/...
nfs_export1  IState             sysa     not waiting
nfs_export1  IState             sysb     not waiting
nfs_export1  MonitorOnly        global   NO
nfs_export1  LastOnline         global
nfs_export1  Enabled            global   YES
nfs_export1  AutoStart          global   YES
nfs_export1  Critical           global   YES
nfs_export1  TriggerEvent       global   NO
nfs_export1  PathName           global   /export1
nfs_export1  Options            global
nfs_export1  OnlineNFSRestart   global   0
nfs_export1  OfflineNFSRestart  global   1
nfs_export1  Type               global   Share
nfs_export1  Group              global   groupx
```

Stopping Groups

To stop both groups and take their resources offline, type the following commands on either system.

```
# hagrp -offline groupx -sys sysa
# hagrp -offline groupy -sys sysb
```

To verify that both groups are offline:

```
# hagrp -display groupx
# hagrp -display groupy
```

Starting Groups

To restart `groupx` on System A and `groupy` on System B, type the following commands on either system:

```
# hagrp -online groupx -sys sysa
# hagrp -online groupy -sys sysb
```

Migrating Groups

To migrate `groupx` from System A to System B, type the following command on either system.

```
# hagrp -switch groupx -to sysb
```

To verify that the group was migrated successfully:

```
# hagrp -display groupx
```

Testing Failover

When both groups are online on System B, test failover by powering down System B. If successful, both groups will fail over to System A.

1. To verify that both groups are online, type the following commands on System A:

   ```
   # hagrp -display groupx
   # hagrp -display groupy
   ```

2. Power up System B and start VCS:

   ```
   # hastart
   ```

3. To migrate `groupy` back, type the following command on either system:

   ```
   # hagrp -switch groupy -to sysb
   ```

Administering VCS from the GUI

The VCS graphical user interface “Cluster Manager” enables you to monitor and administer VCS. Before you can use the Cluster Manager, you must:

✓ Add a user to the configuration to establish a user account.

✓ Set up the GUI display.

Adding A User to the VCS Configuration

To establish a user account for Cluster Manager, execute the following commands as `root` on any system in the cluster.

To establish an account as read/write

1. Set the VCS configuration file to read / write mode:

   ```
   # haconf -makerw
   ```

2. Type the following command to add the user:

   ```
   # hauser -add user_name
   ```

3. Enter a password when prompted.
4. Set the VCS configuration file to read-only:

   # **haconf -dump -makero**

To establish an account as read-only

1. Set the VCS configuration file to read / write mode:

   # **haconf -makerw**

2. Type the following command to add the user:

   # **hauser -add VCSGuest**

3. Press Return when prompted for a password. *Passwords are not required for read-only access.*
4. Set the VCS configuration file to read-only:

   # **haconf -dump -makero**

Setting the Display

The UNIX version of Cluster Manager requires an X-Windows desktop.

To set the display

1. Type the following command to grant the system permission to display on the desktop:

   # **xhost +**

2. Configure the shell environment variable `DISPLAY` on the system where Cluster Manager will be launched. If using the Korn shell, type the following command to display on the system “myws”:

   # **export DISPLAY=myws:0**

Starting Cluster Manager

When you have established a user account and set the display, type the following command to start Cluster Manager:

```
# hagui
```

There are four windows in the Cluster Manager:

- Cluster Monitor, from which you log in, add clusters, and change preferences. This window also provides summary information on each monitored cluster.
- Cluster Explorer, from which you monitor systems and their associated service groups, resources, attributes, and dependencies.
- Log Desk, from which you monitor log messages received from the engine and view commands issued from the GUI.
- Command Center, from which you build VCS commands and send them to the engine.

For more information on using the Cluster Manager, see the chapter on administering VCS from the GUI in the *VERITAS Cluster Server User's Guide*.

Troubleshooting VCS

For more information on troubleshooting VCS, refer to the *VERITAS Cluster Server User's Guide.*

Reconciling Major/Minor Numbers

1. Use the following command to display the major and minor numbers for each of the block devices supporting NFS services from each system. For Volume Manager volumes, you must first import the associated shared disk group on each system, then type the following command:

   # ls -lL *block_device*

   The variable *block_device* refers to the name you identified on page 10.

   For example, type:

   # ls -lL /dev/dsk/c1t1d0s3

   Output from each system resembles:

   On System A:

   ```
   crw-r-----  1 root  sys  32,134 Dec 3 11:50 /dev/dsk/c1t1d0s3
   ```

   On System B:

   ```
   crw-r-----  1 root  sys  36, 62 Dec 3 11:55 /dev/dsk/c1t1d0s3
   ```

2. If the major numbers match, proceed to step 4.

3. If they do not match, complete steps a-f, below. (In the example above, the major numbers are 32 and 36.)

   a. Place the VCS command directory in your path. For example:

      ```
      # export PATH=$PATH:/opt/VRTSvcs/bin
      ```

   b. If the block device is a volume, identify on each system the two major numbers used by the VERITAS Volume Manager:

      ```
      # grep vx /etc/name_to_major
      ```

      Output on System A would resemble:

      ```
      vxio 32
      vxspec 33
      ```

      On System B:

      ```
      vxio 36
      vxspec 37
      ```

   c. Type the following command on System B to change the major number (36/37) to match that of System A (32/33):

      For disk partitions:

      ```
      # haremajor -sd major_number
      ```

      For volumes:

      ```
      # haremajor -vx major_number1 major_number2
      ```

      The variable *major_number* represents the numbers from System A.

      For example, for disk partitions:

      ```
      # haremajor -sd 32
      ```

      For volumes:

      ```
      # haremajor -vx 32 33
      ```

If this command fails, you will receive a report similar to the following:

```
Error: Preexisiting major number 32
These are available numbers on this node: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

d. If you receive this report, type the following command on System A to change the major number (32/33) to match that of System B (36/37):

For disk partitions:

```
# haremajor -sd 36
```

For volumes:

```
# haremajor -vx 36 37
```

If the command fails again, you will receive a report similar to the following:

```
Error: Preexisiting major number 36
These are available numbers on this node: 126...
Check /etc/name_to_major on all systems for
available numbers.
```

e. If you receive the second report, choose the larger of the two available numbers (in this example, 128), and use this number in the haremajor command to reconcile the major numbers. Type the following command on both systems:

For disk partitions:

```
# haremajor -sd 128
```

For volumes:

```
# haremajor -vx 128 129
```

f. Reboot each system on which haremajor was successful.

4. If the minor numbers match, proceed to reconcile the major and minor numbers of your next partition.

5. If the block device on which the minor number does not match is a volume, consult the manual page `vxdg(1M)` for instructions on reconciling the VERITAS Volume Manager minor numbers, with specific reference to the `reminor` option.

6. For disk partitions, complete steps a–e, below. (In this example, the minor numbers are `134` and `62`.)

   a. Type the following command on both systems, using the name of your block device:

      ```
      # ls -1 /dev/dsk/c1t1d0s3
      ```

      Output from this command resembles the following on System A:

      ```
      lrwxrwxrwx  1 root  root  83 Dec 3 11:50 \
        /dev/dsk/c1t1d0s3 -> ../../ \
        devices/sbus@1f,0/QLGC,isp@0,1000/sd@2,0:d,raw
      ```

      The device name (in bold, above) includes the slash following the word `devices`, and continues to, but does not include, the colon.

b. Type the following command on both systems to determine the instance numbers used by the SCSI driver:

```
# grep sd /etc/path_to_inst | sort -n -k 2,2
```

Output from this command resembles the following on System A:

```
"/sbus@1f,0/QLGC,isp@0,10000/sd@0,0" 0 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@1,0" 1 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@2,0" 2 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@3,0" 3 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@4,0" 4 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@5,0" 5 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@6,0" 6 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@8,0" 7 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@9,0" 8 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@a,0" 9 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@b,0" 10 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@c,0" 11 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@d,0" 12 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@e,0" 13 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@f,0" 14 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@0,0" 15 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@1,0" 16 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@2,0" 17 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@3,0" 18 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@4,0" 19 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@5,0" 20 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@6,0" 21 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@8,0" 22 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@9,0" 23 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@a,0" 24 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@b,0" 25 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@c,0" 26 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@d,0" 27 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@e,0" 28 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@f,0" 29 "sd"
```

c. Locate the device names in the output of step b, and identify the instance numbers that appear as the second field in each line. In this example, the device name on System A is `/sbus@1f,0/QLGC,isp@0,10000/sd@1,0`. The associated instance number is `1`.

d. Compare instance numbers.

- If the instance number from one system is not used on the other (that is, it does not appear in the output of step b), edit `/etc/path_to_inst` to make the second system's instance number equal to that of the first system.
- If the instance numbers are being used on both systems, edit `/etc/path_to_inst` on both systems. Change the instance number associated with the device name to an unused number greater than the highest number used by other devices. The output of step b shows the instance numbers used by all devices.

e. Type the following command to reboot each system on which `/etc/path_to_inst` was modified:

```
# reboot -- -rv
```